

DCTNET AND PCANET FOR ACOUSTIC SIGNAL FEATURE EXTRACTION

Yin Xian^{*}, Andrew Thompson[†], Xiaobai Sun[‡], Douglas Nowacek^{*} and Loren Nolte^{*}

^{*} Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

[†] Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

[‡] Department of Computer Science, Duke University, Durham, NC 27708, USA

{yin.xian, xiaobai.sun, doug.nowacek, loren.nolte@duke.edu}, {thompson@maths.ox.ac.uk}

ABSTRACT

We introduce the use of DCTNet, an efficient approximation and alternative to PCANet, for acoustic signal classification. In PCANet, the eigenfunctions of the local sample covariance matrix (PCA) are used as filterbanks for convolution and feature extraction. When the eigenfunctions are well approximated by the Discrete Cosine Transform (DCT) functions, each layer of PCANet and DCTNet is essentially a time-frequency representation. We relate DCTNet to spectral feature representation methods, such as the short time Fourier transform (STFT), spectrogram and linear frequency spectral coefficients (LFSC). Experimental results on whale vocalization data show that DCTNet improves classification rate, demonstrating DCTNet's applicability to signal processing problems such as underwater acoustics.

Index Terms— Convolutional network, PCA, DCT, filterbanks, acoustic perception, spectral clustering, whale vocalizations.

1. INTRODUCTION

The power of multi-layer convolutional networks for learning features has been established in audio signal processing, such as speech recognition and music classification [1, 2, 3]. The idea of a convolutional network is to convolve signals with filters, and use the obtained features for classification. The scattering transform, proposed by Mallat [4], employed pre-specified wavelets as filters, and obtained state of the art result on some music and speech datasets [5, 6]. The structure of the scattering transform is similar to a cascade of constant-Q or mel-filter banks [5], and the scattering transform can capture useful spectral contents of acoustic signals.

The spectral content of acoustic signals is often analyzed using time-frequency energy distributions with a particular frequency scale. For example, the popular mel-frequency cepstral coefficients (MFCC) are based on the mel-scale, a human perceptual frequency scale [7]. The MFCC is a DCT of the log energy of the mel-frequency spectral coefficients (MFSC) [8], which were shown to be closely related to scattering transform coefficients [5, 6].

PCANet [9] was recently proposed for image classification. In this framework, filters are learned from the data as principal components at the local “image patch” level. PCANet was shown to match and in some cases improve upon state of the art performance in a variety of image classification benchmarks.

In this paper, we translate the PCANet framework into the world of acoustic signal processing and relate it to spectral feature representations. The PCANet filters are obtained as the eigenvectors of a local covariance matrix, which is a Toeplitz matrix, and so the resulting filters can be approximated by the Discrete Cosine Transform (DCT) basis functions [10, 11]. We thus introduce the use of DCTNet for acoustic signal classification, in which PCA filters are simply replaced with fixed DCT filters. We relate DCTNet to spectral feature representation methods for acoustic signals, such as the short time Fourier transform (STFT), spectrogram and linear frequency spectral coefficients (LFSC). In particular, each DCTNet layer is essentially a short time DCT. The process of our PCANet and DCTNet is shown in Fig. 1. More technical details can be found in Section 2 to Section 4 and also in the dissertation of Xian [12].




Fig. 1. PCANet and DCTNet Process. The input is the time series of an acoustic signal. After convolving with DCT or PCA filterbanks, we have the short time DCT or short time PCA of the signal. Summing and averaging the second layer's outputs, we have linear scale spectral coefficients, and we use them for spectral clustering and classification.

DCTNet is of interest as an alternative to the scattering transform for the following reasons. Firstly, DCTNet features explicitly give frequency information, which is likely beneficial for many acoustic signals, whereas scattering transform coefficients capture scale information. Secondly, the DCT is a popular tool in audio signal coding [13], for example in MP3, and so DCTNet provides a natural way to incorporate software and hardware in these contexts into a deep learning framework.

We note that Ng and Teoh [14] recently introduced the use of a DCTNet variant for image feature extraction, applying block-wise 2D convolution, histogramming and binary hashing. In contrast, we adopt DCTNet for acoustic signals, using an entirely different post-processing strategy, and we also provide insight into the different time-frequency content revealed by each layer of DCTNet.

We present experiments using the DCLDE whale vocalization data [15]. Results show that DCTNet improved results in classification tasks, suggesting that DCTNet is an attractive tool for acoustic signal processing problem, such as underwater acoustics.

2. THE DCT APPROXIMATION FOR PCANET EIGENFUNCTIONS

In the PCANet framework [9], filterbanks are obtained as the eigenfunctions of the local covariance matrix. Given a signal $\mathbf{x} = (x(1), x(2), \dots, x(N))$, we construct the Hankel matrix

$$\mathbf{X} = \begin{bmatrix} x(1) & x(2) & \cdots & x(N-M+1) \\ x(2) & x(3) & \cdots & x(N-M+2) \\ \vdots & \vdots & \ddots & \vdots \\ x(M) & x(M+1) & \cdots & x(N) \end{bmatrix},$$

where $M < N$. When the first column $x(1), \dots, x(M)$ and the last column $x(N-M+1), \dots, x(N)$ are zeros, letting $\rho_j = \sum_i x(i)x(i+j)$, the sample covariance matrix

$$\mathbf{X}\mathbf{X}^T = \begin{bmatrix} 1 & \rho_1 & \rho_2 & \cdots & \rho_{M-1} \\ \rho_1 & 1 & \rho_1 & \cdots & \rho_{M-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_{M-1} & \rho_{M-2} & \rho_{M-3} & \cdots & 1 \end{bmatrix}$$

is a Toeplitz matrix.

When the autocorrelation of the signal decays fast, the discrete cosine transform (DCT) basis functions can well approximate the eigenfunctions of the Toeplitz matrix [10, 11]. A comparison of the top eight eigenfunctions of DCT and PCA for a single whale vocalization are shown in Fig. 2.

The autocorrelation of the signal and the correlation between the DCT and PCA eigenfunctions are shown in Fig. 3. We can see that the approximation of DCT to PCA eigenfunctions depends upon the structure of the data. An error bound




Fig. 2. Comparison of the top eight DCT eigenfunctions and PCA eigenfunctions for a single whale vocalization

for using DCT eigenfunctions to diagonalize the Toeplitz matrix has been given in terms of the autocorrelation coefficients of the signal [11].




Fig. 3. Signal autocorrelation and eigenfunctions correlation

We can also interpret the DCT approximation in a Toeplitz-Fourier framework, since a Toeplitz matrix T can be represented as the sum of a circulant matrix C and a skew-circulant matrix S , that is $T = C + S$. The eigenfunctions of a circulant matrix are well known to be Fourier eigenfunctions. We can view the skew-circulant matrix as an error, and the energy of the circulant matrix within the given Toeplitz matrix can be optimized by applying Chan's optimal circulant preconditioner [16], that is $C = \arg \min_{C'} \|C' - T\|_F^2$. By this argument, the eigenfunctions are well approximated by Fourier basis functions, and consequently by the closely related DCT basis functions.

3. SHORT TIME PCA AND SHORT TIME DCT

The discrete time STFT can be written as [17]:

$$\begin{aligned} X(m, \omega) &= \sum_{n=-\infty}^{\infty} x(m-n)w(n) \exp(-j\omega(m-n)) \\ &= \exp(-j\omega m) \sum_{n=-\infty}^{\infty} (w(n) \exp(j\omega n))x(m-n) \\ &= \exp(-j\omega m) [(w(m) \exp(j\omega m)) * x(m)] \end{aligned}$$

where w is a window function and ω is the angular frequency. It can be viewed as the modulation of a band-pass filter. For PCANet and DCTNet, we replace the Fourier basis functions $\{\exp(j\omega m)\}$ with PCA eigenfunctions and DCT eigenfunctions, obtaining a short time PCA and short time DCT of the signal respectively.

Plots of short time PCA and DCT (output of the first layer) are shown in Fig. 4 for different window lengths. We use the DCLDE blue whale vocalization data as an example. DCT filterbanks are a natural choice since they are time-frequency representations, whereas time-frequency representation and resolution may be lost when using PCA filterbanks, especially when the PCA eigenfunctions cannot be approximated by the DCT. We can see in the comparison of Fig. 4(a) and Fig. 4(c) that PCA fails to represent the time-frequency content of the signal.

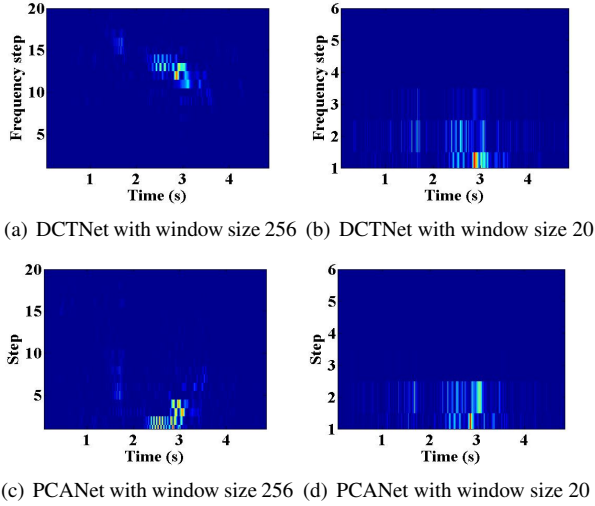


Fig. 4. Plots of the first layer output

The advantage of DCTNet over PCANet can also be found by looking at their computation complexity. PCA is data dependent, and requires $O(M^3)$ operations to find the eigenvectors of the $M \times M$ data covariance matrix, plus an additional computational cost to perform the short time PCA convolution; while for DCTNet, with the help of FFT, we need $O(MN \log_2 M)$ operations to obtain a short time DCT of the signal of length N [14, 18].

4. LINEAR FREQUENCY SPECTROGRAM

After obtaining the first layer, we treat each row of the output as a separate signal, and convolve each with a new PCA or DCT filterbank. We thus end up with multiple new short time PCAs/DCTs, which can capture the dynamic structure of the signal. We choose a smaller window compared with the first layer window size for the filterbanks, so we have a finer scale representation (second layer) inside a coarser scale representation (first layer). Plots of second layer short time DCT are shown in Fig. 5. They are generated by taking a short time DCT of window size 20 of the first layer output signal using window size 256, as shown in Fig. 4(a). Comparing Fig. 4(b) with Fig. 5, we can see that the two-layer DCTNet can illustrate more detailed signal components.




Fig. 5. DCTNet second layer output with window size 256 at the first layer, and window size 20 at the second layer. (a) shows the signal component at frequency step 14 of Fig. 4(a), while (b) shows the signal component at frequency step 12 of Fig. 4(a).

The mel-frequency spectrogram, or MFSC, can be obtained by convolving the signal with a constant-Q filterbank and a low-pass filter, as the analysis in [5, 6]:

$$Mx(t, \lambda) \approx |x * \psi_\lambda|^2 * |\phi|^2(t),$$

where ψ_λ is a constant-Q band-pass filter, and ϕ is a low-pass filter. When replacing the constant-Q filterbank and the low-pass filter with a DCT filterbank and a DCT function, the average of the two-layer DCTNet output can give us a linear frequency spectrogram like features, due to the linearity of scale of DCT filterbank. The average of the two-layer DCTNet features can also improve the deformation stability for acoustic signals.

As Fig 6 shows, we use a piece of classical music from the GTZAN dataset as an example to illustrate MFSC, LFSC, and their connection with the scattering transform and DCTNet outputs. Applying triangular weight functions for the filterbanks of MFSC and LFSC, we obtain 40 MFSC and LFSC coefficients. Most of the energy of this music is distributed in frequency range below 4000Hz. The mel-scale is approximately a linear scale for low frequencies [19], but logarithmic for high frequencies, since the human ear is less sensitive to high frequency signal content [20].

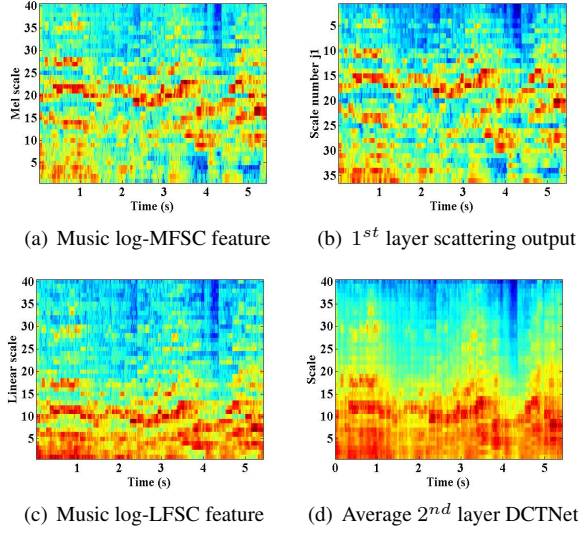


Fig. 6. Comparison of MFSC, scattering transform coefficients, LFSC and the average of DCTNet second layer output

5. EXPERIMENTAL RESULTS

5.1. Dataset

We use the DCLDE 2015 conference blue whale D call data and fin whale 40Hz call data [15] for experiments. There are 851 blue whale calls and 244 fin whale calls of the same sampling frequency 2000Hz. Each signal lasts for 5 seconds, so the length of each signal is 10000.

5.2. Spectral Clustering

We use two-layer DCTNet and PCANet, and use window size 256 for the first layer in order to have a better time-frequency resolution of the data, and choose window size 20 for the second layer to make it comparable to MFSC and LFSC. For MFSC and LFSC, we generate the spectrogram using a Hamming window of length 256. We extract 20 coefficients from each time frame after applying the triangular weighted function filterbank on the signals, and then we concatenate the coefficients along the time axis. We use the MATLAB scattering transform toolbox [21] with the frame duration $T = 125ms$ (or window size 256) for the experiment, and set $Q_1 = 12$, $Q_2 = 1$ to better capture the acoustic signal information. In this experiment, the energy of the second layer scattering transform coefficients is very close to zero. Considering the size of dataset and features characteristic, to make a fair experiment, we compared with one layer scattering transform coefficients, MFSC, and LFSC.

We use the Laplacian Eigenmap for dimensional reduction for the feature set of DCTNet, PCANet, scattering transform, MFSC and LFSC. We use three nearest neighbors to

create the kernel of the Laplacian Eigenmap. We examine the adjacency matrices created by the Laplacian Eigenmap, which are generated based on feature similarity. Since it is a binary classification, we use the Fielder value [22] for spectral re-ordering. We can see that the two-layer DCTNet and PCANet can better separate the blue whale and fin whale data than one layer, as Fig 7 shows.




Fig. 7. Adjacency matrices created by DCTNet and PCANet. The block of rows and columns from 1 to 851 are blue whale data, from 852 to 1095 are fin whale data.

After obtaining features from DCTNet, PCANet, the scattering transform, MFSC and LFSC, we use the kNN classifier ($k = 3$) to evaluate the separation of the data. The ROC plot is shown in Fig. 8. The values of the AUC (area under the ROC curve) are shown in Table 1.

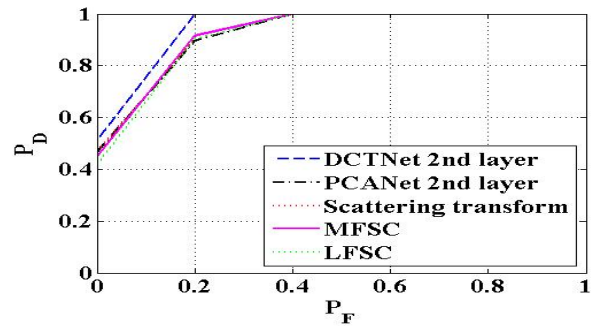


Fig. 8. ROC comparisons

We use 5-fold cross-validation to generate the plot, that is 80% of the blue whale vocalizations and fin whale vocalizations for training, and the rest for testing. We calculate the distance of each testing data point to the training data points, and make a decision based on its three nearest neighbors, and compare it with the ground truth, to obtain the true positive rate and false positive rate. In order to generate the ROC

Table 1. Area Under the Curves (AUCs) of different feature sets

Feature set	AUC
DCTNet 2 nd layer	0.9513
PCANet 2 nd layer	0.9258
DCTNet 1 st layer	0.9200
PCANet 1 st layer	0.9079
Scattering transform	0.9275
MFSC	0.9283
LFSC	0.9225

curve, we vary the value of the probability of false alarm (P_F) over the range 0 to 1, and obtain a corresponding threshold for the probability of detection (P_D) based upon the obtained true positive and false positive rate. Since the number of blue right whale and fin whale vocalizations available for testing are relatively small, the classification results are promising but preliminary.

6. CONCLUSION

In this paper, we apply PCANet and DCTNet for acoustic signal classification. We have shown that each layer of PCANet and DCTNet is essentially a short time PCA/DCT, revealing different time-frequency content of the signal. The computation cost for DCTNet is smaller than for PCANet. DCTNet can generate linear frequency spectral like coefficients, and improve the classification rate of whale vocalizations, which is useful for underwater acoustics, or potentially other applications.

7. REFERENCES

- [1] Y. LeCun and Y. Bengio, “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, 1995.
- [2] G. Hinton, L. Deng, D. Yu, G. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, and T. Sainath, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *Signal Processing Magazine, IEEE*, vol. 29, no. 6, pp. 82–97, 2012.
- [3] H. Lee, P. Pham, Y. Largman, and A. Ng, “Unsupervised feature learning for audio classification using convolutional deep belief networks,” in *Advances in neural information processing systems*, 2009, pp. 1096–1104.
- [4] S. Mallat, “Group invariant scattering,” *Communications on Pure and Applied Mathematics*, vol. 65, no. 10, pp. 1331–1398, 2012.
- [5] J. Andén and S. Mallat, “Multiscale scattering for audio classification,” in *ISMIR*, 2011, pp. 657–662.
- [6] J. Andén and S. Mallat, “Deep scattering spectrum,” *Signal Processing, IEEE Transactions on*, vol. 62, no. 16, pp. 4114–4128, 2014.
- [7] S. Stevens, J. Volkman, and E. Newman, “A scale for the measurement of the psychological magnitude pitch,” *The Journal of the Acoustical Society of America*, vol. 8, no. 3, pp. 185–190, 1937.
- [8] L. Rabiner and B. Juang, “Fundamentals of speech recognition,” 1993.
- [9] T. Chan, K. Jia, S. Gao, J. Lu, Z. Zeng, and Y. Ma, “PCANet: A simple deep learning baseline for image classification?,” *arXiv preprint arXiv:1404.3606*, 2014.
- [10] N. Ahmed, T. Natarajan, and K. Rao, “Discrete cosine transform,” *Computers, IEEE Transactions on*, vol. 100, no. 1, pp. 90–93, 1974.
- [11] V. Sánchez, P. Garcia, A. Peinado, J. Segura, and A. Rubio, “Diagonalizing properties of the discrete cosine transforms,” *Signal Processing, IEEE Transactions on*, vol. 43, no. 11, pp. 2631–2641, 1995.
- [12] Y Xian, *Detection and classification of whale vocalizations*, Ph.D. thesis, Duke University, 2015.
- [13] J. Princen and A. Bradley, “Analysis/synthesis filter bank design based on time domain aliasing cancellation,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 34, no. 5, pp. 1153–1161, 1986.
- [14] C. Ng and A. Teoh, “DCTNet: A simple learning-free approach for face recognition,” *arXiv preprint arXiv:1507.02049*, 2015.
- [15] Scripps Institution of Oceanography, “DCLDE conference low frequency data,” <http://www.cetus.ucsd.edu/dclde/dataset.html>, 2015.
- [16] T. Chan, “An optimal circulant preconditioner for toeplitz systems,” *SIAM journal on scientific and statistical computing*, vol. 9, no. 4, pp. 766–771, 1988.
- [17] A. Oppenheim, *Applications of digital signal processing*, Prentice-Hall signal processing series. Prentice-Hall, 1978.
- [18] R. Wang, *Introduction to orthogonal transforms: with applications in data processing and analysis*, Cambridge University Press, 2012.
- [19] B. Logan, “Mel frequency cepstral coefficients for music modeling,” in *ISMIR*, 2000.

- [20] H. Olson, *Music, physics and engineering*, vol. 1769, Courier Corporation, 1967.
- [21] J. Anden and S. Mallat, “Scatnet (v0.2),” <http://www.di.ens.fr/data/software/scatnet/>, 2013.
- [22] U. Von Luxburg, “A tutorial on spectral clustering,” *Statistics and computing*, vol. 17, no. 4, pp. 395–416, 2007.